Quantifying Benefits of a Personal Knowledge Assistant on Task Resumption

Christian JILEK[1,2], Paula GAUSELMANN[3], Jessica CHWALEK[1,2], Tobias TEMPEL[3], Andreas DENGEL[1,2]

[1] *German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany*
[2] *TU Kaiserslautern, Kaiserslautern, Germany*
[3] *Ludwigsburg University of Education, Ludwigsburg, Germany*

Abstract. Due to constant information overload in knowledge work we are working towards a self-(re-)organizing, especially forgetful, context-sensitive assistance system (cSpaces). Here, we present a study on task (context) switching in which participants were supported by an early version of our system. Participants were assigned to either a control group with no, a group with context-insensitive or a group with context-sensitive support by the assistance system. Activities were logged and analyzed regarding group-dependent differences in task resumption lags. Participants supported by our assistant took less time to resume a former activity after context switching than participants with no support. This result provides initial evidence that context-sensitive, intelligent systems like cSpaces are able to effectively support knowledge workers.

Keywords: knowledge work, context-sensitive assistance and user interfaces, user activity tracking, multi-tasking, task resumption

1. Introduction

Knowledge workers face an ever increasing flood of information in their daily work, which involves activities like finding and organizing, processing, assessing, and creating information. Such information overload can lead to a decrease in cognitive resources (Miller, 1956; Klingberg, 2009). Cognitively offloading part of one's mental load, e.g. through saving, can counteract this overload and lead to better performance in currently relevant tasks (Runge et al., 2019). However, the norm is still that the individual is responsible for his or her information management, i.e. for deciding what is stored when, where and how. Frequently switching from one work context (e.g. designing an experiment) to another (e.g. correcting student's papers), each being associated with different tasks, documents, mails, etc., makes personal information management a challenging and often impossible task. As a typical consequence, users' personal information sphere, such as mails, files, bookmarks and folder hierarchies, is cluttered from being used in many different contexts. They typically do not find the time to rehash or tidy up. Regardless of problematic information management, there is evidence that mental access to context-specific information is impaired after context-change (Sahakyan & Kelley, 2002). Information chaos makes this access even more difficult as cognitive resources have to be used solely to filter irrelevant information and to mentally reactivate and find content of a currently relevant work context after having worked in another, unrelated one.

To counter this we have developed a self-(re)organizing (and especially forgetful), context-sensitive knowledge assistance system called *Context Spaces* (Jilek et al.,

2018) - or *cSpaces* for short. The name originates from two of its core ideas: first, the capturing of user contexts as if the system was a butler standing next to the user while working on their computer. This is necessary to learn enough about the user's world in order to decide on and execute appropriate support measures. In this way, a machine understandable, knowledge graph-based representation of the user's mental model, a *personal information model* (Sauermann et al., 2007), is captured and utilized. Second, we see context as an active and explicit element users can work with. As support measures we particularly think of *Managed Forgetting* (Jilek et al., 2019), by which we understand an escalating set of measures overcoming the binary keep-or-delete paradigm. They range from temporal hiding, to condensation and reorganization, to adaptive archiving, synchronization and deletion. Several of the system's features need some time to have passed, which is why they are evaluated in long-term studies. Here, we focus primarily on context capturing and task switching aspects.

2. Method

Participants: The study involved 51 participants (33 female) with an average age of 26.3, mostly students of humanities.

Material: All participants worked with a prototype of *cSpaces* specifically adapted to the experiment. This prototype came with a context-sensitive sidebar consisting of different widgets that are shown and explained in Figure 1.

40 fictive, standardized websites were used for the study (seven mails, 33 files and websites). It shall be noted that, for the sake of a more controlled experiment, files and mails used for the experiment were moved to the browser, thus being available as if using a respective online/cloud service. As a consequence, logging each action was reliable and granting or denying access to each information item could be fully controlled.

Procedure: Participants had to plan a barbecue evening which involved getting an overview what people would like to eat and drink, what was brought to the party by guests and what still had to be ordered. Additionally, they had to check various vendors for prices and offerings to finally fill an order list containing everything to be ordered for the party. For this, they had to read files, mails and websites. Participants were split into three groups: G1 as a control group only had the content analysis widget (see 1 in Fig. 1), i.e. a non-interactive sidebar. G2 had the additional possibility to add items to a single context-free folder (3) as well as adding notes to them.

Their activities were also captured in a history list (4). G3 also had these possibilities but with the additional help of items and activities being separated by contexts (2). G3 thus was the only group being supported by all four widgets depicted in Fig. 1. After a short tutorial on using the prototype, users began working on the main task. They were given the information that they would be interrupted to switch to a second task - from the perspective of the experimenter: a distraction task – which was about working on a wiki entry about autonomous driving. Each participant spent in total 20 minutes on the main and 12 minutes on the distraction task. Three switches between the two were triggered leading to an overall number of 153 observed context switches as well as one hour of video material and 1000 logged activities on average for each participant. Quantifying *task resumption lags* (TRL) (Altmann & Trafton, 2004), by which we understand the time a person needs to readjust their mind to the former context after switching back, was performed by measuring the time until a participant got back to a page actually containing content (i.e. offerings or amounts).

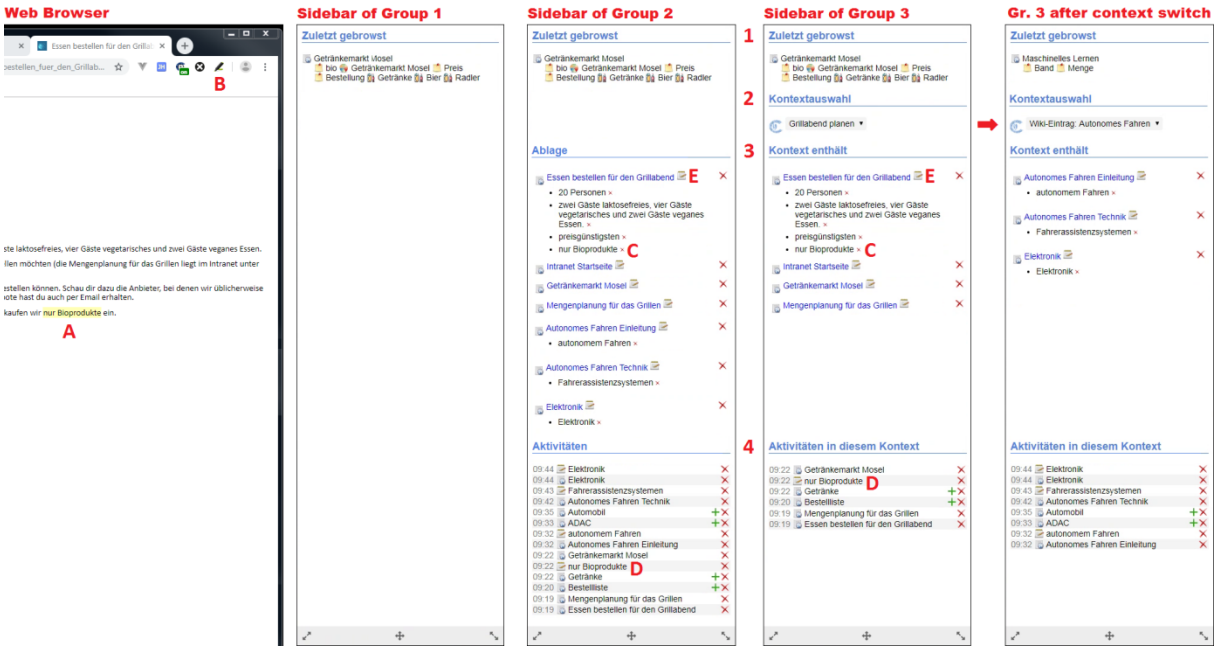

Figure 1. *In the experiment, participants browsed websites and the cSpaces sidebar reacted accordingly. Each of the three groups had a different combination of sidebar widgets as depicted in the middle of the figure. G1 only had a small content analysis widget: cSpaces performed information extraction on the website's content and showed recognized entities (1). One of the most important inputs was the context selection (2) right below. Users of G3 could select in which context they wished to work. Selecting (or switching) the current context reorganized the sidebar, which, in the experiment, consisted of a list of information items associated with the context ("context contains", 3) and a history of actions performed in this context ("activities in this context", 4), e.g. viewed information elements, notes taken, etc. Participants of G2 also had (3) and (4) but without the possibility to switch between contexts. They just had a single context for both tasks, thus items and activities were mixed up. Users could add the information item associated with an activity (e.g. a browsed website) to the context by clicking the green plus next to the activity's description. Removing activities from the list could be done with the red cross next to it. Clicking an item in the context opened it in the web browser. For each item users could take additional notes: clicking the notepad icon (E) opened a text area below the respective item. Again, there was a red cross to delete a note. There was also the possibility to select text passages (A) and click the highlighting button (B) developed as an auxiliary tool to the sidebar. This highlighted the respective section in the browser and created a read-only note associated with the current website (C). If the website had not been added to the context before, this was done automatically before storing the note. Additionally, the action of taking a note was added to the context's activity history (D). The right-hand side of the figure shows participants of G3 switching their context from planning a barbecue evening to working on a wiki text on autonomous driving: the sidebar content was updated accordingly.*

For this purpose, we classified all 40 pages into offerings and prices (17), amounts of food and beverages (5), navigation pages (10), spam (7), or the order list (1). Using the captured video material all measurements were verified and corrected if necessary. For example, if a user navigated through several pages in quick succession (<1sec), we assumed that the detention time was too short to actually conceive anything written on that page. We thus used the next content page visit or sidebar addition for the TRL calculation. Note that data sets of four participants belonging to G3 had to be excluded since they did not use the sidebar to do the tasks, i.e. they had zero or one sidebar actions during the whole experiment.

3. Results

We conducted unpaired t-tests of G1 vs. G2 and G1 vs. G3: n_1=63, n_2=27, n_3=51, $\bar{x}_1$=18273.14ms, $\bar{x}_2$=10969.59ms, $\bar{x}_3$=10548.14ms, $\bar{s}_1$=8196.19ms, $\bar{s}_2$=5770.1ms, $\bar{s}_3$=6741.81ms and additional u-tests having p<0.001. The average TRL for group G3 having a context-sensitive sidebar was 10.5 seconds, followed by the TRL of G2 of 11.0 sec, and finally a value of 18.3 sec for the control group G1. Comparing G2 and G3 to the control group G1 yields Hedges' g (Hedges, 1981) of 0.97 and 1.02, respectively, indicating a large effect in both cases. Results are significant (p<0.001).

4. Discussion

In the present data study we examined how and to what extent our intelligent, context-sensitive assistance system *cSpaces* affects task resumption lags after repeated context-switching. Participants supported by an interactive sidebar (G2, G3) took significantly less time to resume a former task after context-switching than participants with no support (G1).

Previous studies showed that cognitively offloading encoded content via saving has a beneficial effect on subsequent task performance (Storm & Stone, 2015; Runge et al., 2019). In contrast to participants from G2 and G3, whose activities were constantly saved, updated and displayed in an interactive sidebar, participants from G1 were forced to mentally maintain currently not relevant content after context switching. This extra strain on mental resources can be one possible reason for delayed TRLs in G1. Additionally, memory is context-sensitive. A change of context, as was the case here several times, impairs mental access to previously encoded content (Sahakyan & Kelley, 2002).

Another approach to explaining the difference between subjects with and without support by our system is based on the so-called "two faces" of memory retrieval. Depending on the circumstances, retrieval cues can either improve or impair subsequent retrieval, that is, recalling further information in addition to the cued information itself. Goenert & Larson (1994), for example, showed that the effect of providing a subset of previously encoded information varied depending on whether subjects were instructed to either remember or to forget a memorized item set before learning a second set. Afterwards, they were tested on the first set, with part of its items being present as retrieval cues. For participants who were told to remember the first set, those cues led to impaired recall of the remaining items. In contrast, participants who were told to forget the first set showed better recall performance with the retrieval cues being present during test.

The contents of the sidebar in the form of notes, headlines, etc. which were available for G2 and G3 can be seen as similar retrieval cues. Participants in these groups were allowed to forget the content of a currently not relevant task as it was constantly saved in the sidebar. When returning to their previous task, they most likely used this content to help them mentally find back into it. Based on the results of the study described above, shorter TRLs in G2 and G3 can be interpreted as a consequence of retrieval-cue facilitation.

The contents of the sidebar used in this study can be seen as similar retrieval cues.

An additional factor that influences the effects of memory retrieval is context. In a study by Bäuml & Samenieh (2012), participants, after having memorized a first word list, were asked to either perform a neutral counting task or (mentally) change their

internal context before having to memorize a second list. Afterwards, they were tested for list one with part of its items being present as retrieval cues. Participants who did not change their internal context showed retrieval-cue impairment, i.e. impaired recall performance when retrieval cues were present at test. In contrast, participants who did change their internal context showed facilitation, i.e. improved recall performance when using retrieval cues at test. Thus, the same effect occurred as after intentional forgetting in the study by Goenert and Larson (1994), which probably also created a mental context change (cf. Sahakyan & Kelley, 2002).

In the present study, switching between tasks means not only a structural but also an internal context switch. Accordingly, all participants experienced several (internal) context changes. However, while participants from G1 had no information available to help them return to a former context, participants who were supported by an interactive, self-updating sidebar, which could thus have led to shorter TRLs.

Unexpectedly, however, there was no significant difference between participants from G2 (context-insensitive sidebar) and G3 (context-sensitive sidebar). We would have expected that a context-specific order of content listed in the sidebar would lead to faster task resumption than if contents of different contexts were mixed. It is possible that no such difference could be found as, here, participants had to merely switch between two different tasks, one of which was simply a distraction task and discovered as such by some participants. In everyday work it is common to work in numerous different contexts, switching from one to the next with new ones constantly accumulating. We will further investigate this in a long-term study already in process.

5. Conclusion & Outlook

Although the long-term evaluation of a more feature-rich knowledge assistant is still in progress, we could already observe interesting effects when re-analyzing captured data of studies conducted so far. In contrast to former studies on web-based multitasking, we explicitly quantified user's task resumption lag when switching contexts while being supported by an AI-based agent in different ways. The difference between the most supported group in the test and the control group was about 42%, meaning they could resume tasks almost twice as fast. The difference between both supported groups was insignificant, presumably because of too few contexts, which we now address in a long-term study. In addition, more sophisticated support measures like hiding or reorganization are available in that study.

6. References

Altmann, E. M. & Trafton, J. G. (2004) Task interruption: Resumption lag and the role of cues. In: Proc. of the Annual Meeting of the Cognitive Science Society. Vol. 26. 26.

Bäuml, K. H. T. & Samenieh, A. (2012). Selective memory retrieval can impair and improve retrieval of other memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(2), 488.

Goernert, P. N. & Larson, M. E. (1994). The initiation and release of retrieval inhibition. *The Journal of general psychology*, *121*(1), 61-66.

Hedges L. V. (1981) Distribution Theory for Glass's Estimator of Effect size and Related Estimators. In:Journal of Educational Statistics 6.2 (1981), pp. 107–128.

Jilek, C., Schröder, M., Schwarz, S., Maus, H. & Dengel, A. (2018) Context Spaces as the Cornerstone of a Near-Transparent and Self-Reorganizing Semantic Desktop. In: The Semantic Web: ESWC 2018 Satellite Events, Revised Selected Papers. Springer, 2018, pp. 89–94.

Jilek, C., Runge, Y., Niederée, C., Maus, H., Tempel, T., Dengel, A. & Frings, C. (2019) Managed Forgetting to Support Information Management and Knowledge Work. In: KI – Künstliche Intelligenz, German Journal on Artificial Intelligence 33.1 (Mar. 2019), pp. 45–55.

Klingberg, T. (2009). *The overflowing brain: Information overload and the limits of working memory*. Oxford University Press.

Miller, G. (1956). Human memory and the storage of information. *IRE Transactions on Information Theory*, *2*(3), 129-137.

Runge, Y., Frings, C. & Tempel, T. (2019). Saving-enhanced performance: Saving items after study boosts performance in subsequent cognitively demanding tasks. *Memory*, *27*(10), 1462-1467.

Sahakyan, L. & Kelley, C. M. (2002). A contextual change account of the directed forgetting effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(6), 1064.

Sauermann, L., van Elst, L. & Dengel, A. (2007) PIMO – a Framework for Representing Personal Information Models. In: Proc. of I-Media ’07 & I-Semantics ’07. Know-Center, Austria, pp. 270–277.

Storm, B. C. & Stone, S. M. (2015). Saving-enhanced memory: The benefits of saving on the learning and remembering of new information. *Psychological Science*, *26*(2), 182-188.

Acknowledgement. This work was partly funded by the Deutsche Forschungs-gemeinschaft (DFG; DE 420/19-1) and the Federal Ministry of Education and Research (project SensAI; grant no. 01IW20007). We also thank Yannick Runge, Christian Frings (University of Trier), Janina Schuhmacher and Nele Russwinkel (TU Berlin) for their contributions to making the initial study and data collection possible.

Gesellschaft für
Arbeitswissenschaft e.V.

Arbeit HUMAINE gestalten

67. Kongress der
Gesellschaft für Arbeitswissenschaft

Lehrstuhl Wirtschaftspsychologie (WiPs)
Ruhr-Universität Bochum

Institut für Arbeitswissenschaft (IAW)
Ruhr-Universität Bochum

3. - 5. März 2021

GfA-Press

Bericht zum 67. Arbeitswissenschaftlichen Kongress vom 3. - 5. März 2021

Lehrstuhl Wirtschaftspsychologie, Ruhr-Universität Bochum
Institut für Arbeitswissenschaft, Ruhr-Universität Bochum

Herausgegeben von der Gesellschaft für Arbeitswissenschaft e.V.
Dortmund: GfA-Press, 2021
ISBN 978-3-936804-29-4

NE: Gesellschaft für Arbeitswissenschaft: Jahresdokumentation

Als Manuskript zusammengestellt. Diese Jahresdokumentation ist nur in der Geschäftsstelle erhältlich.

Schriftleitung: Matthias Jäger
im Auftrag der Gesellschaft für Arbeitswissenschaft e.V.

Screen design und Umsetzung
© 2021 fröse multimedia, Frank Fröse
office@internetkundenservice.de · www.internetkundenservice.de